

511-312  
N95-18978

1994

NASA/ASEE SUMMER FACULTY FELLOWSHIP PROGRAM

MARSHALL SPACE FLIGHT CENTER  
THE UNIVERSITY OF ALABAMA

PERFORMANCE EVALUATION OF THE  
ENGINEERING ANALYSIS AND DATA SYSTEM (EADS) II

P. 6  
correct #  
assigned  
#30901  
R41 4/20/95

Prepared By: Linda S. DeBrunner, Ph.D.

Academic Rank: Assistant Professor

Institution and Department: The University of Oklahoma  
School of Electrical Engineering

NASA/MSFC:

Laboratory: Information Systems Office  
Division: Systems Development and Implementation Division  
Branch: Computer Systems Branch

MSFC Colleagues: Ellen Williams  
Amy Epps



## INTRODUCTION

The Engineering Analysis and Data System (EADS) II [1] was installed in March 1993 to provide high performance computing for science and engineering at Marshall Space Flight Center (MSFC). EADS II increased the computing capabilities over the existing EADS facility in the areas of throughput and mass storage. EADS II includes a Vector Processor Compute System (VPCS), a Virtual Memory Compute System (VMCS), a Common File System (CFS), a Common Output System (COS), as well as Image Processing Stations, Mini Super Computers, and Intelligent Workstations. These facilities are interconnected by a sophisticated network system. This work considers only the performance of the VPCS and the CFS. The VPCS is a Cray YMP. The CFS is implemented on an RS 6000 using the UniTree Mass Storage System.

To better meet the science and engineering computing requirements, EADS II must be monitored, its performance analyzed, and appropriate modifications for performance improvement made. Implementing this approach requires tool(s) to assist in performance monitoring and analysis. In Spring 1994, PerfStat 2.0 was purchased to meet these needs for the VPCS and the CFS. PerfStat[2] is a set of tools that can be used to analyze both historical and real-time performance data. Its flexible design allows significant user customization. The user identifies what data is collected, how it is classified, and how it is displayed for evaluation. Both graphical and tabular displays are supported.

We evaluated the capability of the PerfStat tool, suggested and implemented appropriate modifications to EADS II to optimize throughput and enhance productivity, and observed the effects of the modifications on system performance. In this paper, we briefly describe the PerfStat tool, then outline its use with EADS II. Next, we describe the evaluation of the VPCS, as well as modifications made to the system. Finally, we draw conclusions and outline recommendations for future work.

## THE PERFSTAT PERFORMANCE ANALYSIS TOOL

The software architecture of PerfStat is shown in Figure 1. Data collectors are run on the Systems Under Study to collect metrics. More than one data collector may be run on a particular System Under Study (SUS). This data is collected directly from operating system counters and is extracted from system logs.

The data is transferred from each SUS to the Archive Workstation where it is stored. The specific data to be collected from a SUS is determined by a metrics pool and a selection list. The metrics pool defines which metrics can be collected by data collectors on the SUS. The selection list for a particular data collector identifies which of the metrics in the metrics pool will be collected.

After or during data collection, User Interface Workstations may be used to display the data either graphically or in a tabular format. The User Interface Workstation may be located on the same physical computer as the Archive Workstation. Many people may display data simultaneously through the use of multiple user interfaces. The user may classify data in many ways. For example, data may be classified by user id or by level of memory usage. A more

complicated classification might combine user id and level of memory usage to identify processes associated with a particular user id that require more than 4 Mwords of memory.

### USE OF PERFSTAT WITH EADS II

PerfStat is used with EADS II to study two systems: the VPCS and the CFS. A single Archive Workstation is implemented on a separate workstation. The workstation used as the Archive Workstation also serves as a User Interface Workstation. In addition, several other User Interface Workstations are being used.

The Common File System is implemented by a Maximum Strategy Disk Array and robotic tape storage systems controlled by an RS 6000 workstation. UniTree file management software is used. To monitor the CFS, a customized metrics pool that extracts data from the UniTree log files, as well as collecting data from the operating system was provided.

The VPCS is a Cray-YMP supercomputer running UNICOS 6.1. We customized the metrics pool provided to collect information of particular interest. The data is derived from operating system log files and counters.

Several graphs were developed to regularly monitor the CFS and the VPCS. Additional graphs are easily created to investigate new situations.

### EVALUATION AND TUNING OF THE VPCS

After observing the VPCS, several areas were identified for possible improvement. Of these areas, two were selected for implementation during the period of this work. Coincidentally, a previously planned hardware change was implemented to increase system capacity. On July 14, 2 CPUs were added to increase the number of CPUs from 6 to 8. On July 20, several changes were made to the ldcache structure. Then, on July 27, the structure of the Network Queuing System was changed for weekends.

As shown in Figure 2, the addition of CPUs appears to have resulted in a decrease in CPU utilization, as expected. During the same time period, CPU utilization decreased because of a significant workload decrease related to contractor layoffs. This workload decrease is expected to be temporary as the work is shifted to other personnel. In addition, users from the Cray-XMP are being shifted to the Cray-YMP. Weekends are easily identifiable on the graph in Figure 2.

Figure 3 shows the memory requirements for the Cray-YMP. Again, weekends appear as decreases in memory requirements. This graph shows that the current memory of 128 Mwords is adequate for the current situation since the system is not significantly over-subscribed.

The ldcache structure was changed for the /wrk file system. The big file allocation size was set to match the block size. This change should reduce fragmentation[3]. In addition, we eliminated the ldcache for the /bin file system since its re-use ratio was low (analyzed using PerfStat). Third, we increased the number of hash table entries to improve the performance of the ldcache. The effects of these changes were not expected to be evident for several weeks. They will continue to be monitored. Figure 4 shows the re-use ratio for the /wrk file system. In addition, a script was

prepared to measure the disk fragmentation indirectly by allocating a file on each partition of the /wrk file system.

The increase in system idle time, particularly on weekends, led us to investigate the Network Queuing System (NQS) structure. Using PerfStat, we observed jobs ready to run on the weekends, even though the system was taking idle time. To improve system throughput, we increased user limits and queue limits for the weekends. These changes affect system utilization (Figure 1) and queue backlog (Figure 5). PerfStat can be easily used to identify appropriate modifications to the NQS structure as the system workload changes.

### CONCLUSIONS AND RECOMMENDATIONS

PerfStat has proven to be a valuable tool for collecting and analyzing system performance data. Extensive use of the tool with the VPCS has shown that it can be used to identify areas of performance problems, as well as performing routine monitoring and resource forecasting. Additional investigation is needed to determine how this tool can be most effectively used on the CFS.

Several areas are recommended for future modification on the VPCS. When the operating system is upgraded to UNICOS 7.0 in late August 1994, PerfStat can be used to analyze the effects of performance tuning as outlined by UNICOS guidelines [4]. In addition, the NQS performance should be monitored as the workload fluctuates. Changes should be considered for night and "lunch time" NQS structures to compensate for the predictable varying of the workload. Last, the goals of the system should be considered (batch versus interactive, etc.), and appropriate modifications be made to the scheduling parameters.

### REFERENCES

1. *Engineering Analysis and Data System II, Request for Proposal*, RFP No. 8-1-9-AI-00120, George C. Marshall Space Flight Center, NASA, June 19, 1991.
2. *PerfStat Version 2.0 User Guide*, Instrumental, Minneapolis, MN, 1993.
3. *UNICOS Tuning Guide*, SN-2099 6.1, Cray Research, Inc., 1991.
4. *UNICOS Tuning Guide*, SR-2099 7.0, Cray Research, Inc., 1993.

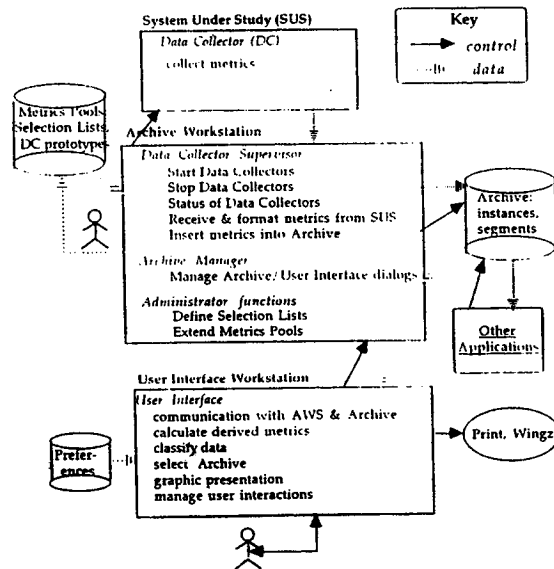


Figure 1. Software Architecture of PerfStat 2.0. [2]




Figure 2. VPCS System Utilization.




Figure 3. VPCS Memory Requirements.




Figure 4. Re-use ratio for ldcache (/wrk).




Figure 5. NQS Backlog